Agile Inexact Methods for Spectral Projector-Based Graph Fourier Transforms

Joya A. Deri, *Member, IEEE,* and José M. F. Moura, *Fellow, IEEE*

arXiv:1701.02851v1 [cs.SI] 11 Jan 2017

***Abstract*—We propose an inexact method for the graph Fourier transform of a graph signal, as defined by the signal decomposition over the Jordan subspaces of the graph adjacency matrix. This method projects the signal over the generalized eigenspaces of the adjacency matrix, which accelerates the transform computation over large, sparse, and directed adjacency matrices. The trade-off between execution time and fidelity to the original graph structure is discussed. In addition, properties such as a generalized Parseval's identity and total variation ordering of the generalized eigenspaces are discussed.**

The method is applied to 2010-2013 NYC taxi trip data to identify traffic hotspots on the Manhattan grid. Our results show that identical highly expressed geolocations can be identified with the inexact method and the method based on eigenvector projections, while reducing computation time by a factor of 26,000 and reducing energy dispersal among the spectral components corresponding to the multiple zero eigenvalue.

***Index Terms*—Jordan decomposition, generalized eigenspaces, directed graphs, sparse matrices, signal processing on graphs, large networks, inexact graph Fourier transforms**

I. Introduction

Graph signal processing [1], [2], [3], [4], [5] allows for analysis of signals over graph structures with the framework of digital signal processing. Properties that have recently been explored include graph filtering [3], [5] and sampling and recovery of graph signals [6], [7], [8]. This paper focuses on considerations for applying adjacency matrix-based graph signal processing as in [1], [9], [10], [11] to real-world networks.

Real-world networks are often characterized by two salient features: (1) extremely large data sets and (2) non-ideal network properties that require extensive computation to obtain suitable descriptions of the network. These features can result in extremely long execution times for the analyses of interest. This paper presents methods that significantly reduce these computation times.

The adjacency matrices of real-world large, directed, and sparse networks may be *defective*. For such matrices, [1] defines the eigendecomposition of the matrix in terms of the Jordan decomposition. This requires determination of Jordan chains, which are expensive to compute, potentially numerically unstable, and non-unique with respect to a set of fixed eigenvectors, resulting in non-unique coordinate graph Fourier transforms. In [11], we reformulate the graph signal processing framework so that the graph Fourier transform is unique over defective networks.

Since our objective is to provide methods that can be applied to real-world systems, it is necessary to consider the associated computational cost. Efficient algorithms and fast computation methods are essential, and with modern computing systems, parallelization and vectorization of software allow decreased computation times [12]. In this paper, we provide a method to accelerate computation by relaxing the determination of the Jordan decomposition of an adjacency matrix.

Consider a graph $\mathcal{G} = G(A)$ with adjacency matrix $A \in \mathbb{C}^{N\times N}$ and a graph signal $s \in \mathbb{C}^N$, which characterizes properties or behaviors at each node of the graph. The graph Fourier transform of [1] is defined as follows. If the adjacency matrix has Jordan decomposition $A = VJV^{-1}$, the graph Fourier transform of a signal $s \in \mathbb{C}^N$ over $\mathcal{G}$ is defined as

$$\widetilde{s} = V^{-1}s, \tag{1}$$

where V^{-1} is the *Fourier transform matrix* of $\mathcal{G}$. This is essentially a projection of the signal onto the proper and generalized eigenvectors of A. It is an orthogonal projection when A is normal ($A^HA = AA^H$) and the eigenvectors form a unitary basis (i.e., $V^{-1} = V^H$).

For *defective*, or non-diagonalizable adjacency matrices, which have at least one eigenvalue λ_i with algebraic multiplicity (the exponent in the characteristic polynomial of A) greater than the geometric multiplicity (the kernel dimension of $A - \lambda_i I$, or $\dim \mathrm{Ker}(A - \lambda_i I)$, where I is the $N \times N$ identity matrix), the corresponding eigenvectors may not span $\mathbb{C}^N$. The basis can be completed by computing Jordan chains of generalized eigenvectors [13], [14], but the computation introduces degrees of freedom that render these generalized eigenvectors non-unique. This is addressed in [11] by defining a spectral projector-based GFT in terms of projections onto the Jordan subspaces of the adjacency matrix of

This work was partially supported by NSF grants CCF-1011903 and CCF-1513936 and an SYS-CMU grant.

The authors are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA (email: jderi,moura@andrew.cmu.edu)

the graph. Formally, the GFT of a signal $s \in \mathbb{C}^N$ over graph $\mathcal{G} = G(A)$, where adjacency matrix A with k distinct eigenvalues has Jordan decomposition $A = VJV^{-1}$ and Jordan subspaces $\{\mathscr{J}_{ij}\}$, $i = 1, \ldots, k$, $j = 1, \ldots, \dim \mathrm{Ker}(A - \lambda_i I)$, is defined as the mapping

$$\begin{aligned} \mathcal{F} : \mathcal{S} &\to \bigoplus_{i=1}^{k} \bigoplus_{j=1}^{g_i} \mathscr{J}_{ij} \\ s &\to (\widehat{s}_{11}, \ldots, \widehat{s}_{1g_1}, \ldots, \widehat{s}_{k1}, \ldots, \widehat{s}_{kg_k}) , \end{aligned} \quad (2)$$

where each $\widehat{s}_{ij}$ represents the projection onto Jordan subspace $\mathscr{J}_{ij}$ parallel to $\mathcal{S} \backslash \mathscr{J}_{ij}$.

The spectral projector GFT requires Jordan chain vector computations for non-diagonalizable, or *defective* adjacency matrices, which can be exceedingly inefficient for large, directed matrices with many nontrivial Jordan subspaces. The desire for a more efficient transform computation motivates the inexact approaches proposed in this paper. In particular, we propose an inexact method for computing the graph Fourier transform that simplifies the choice of projection subspaces. This method is motivated by insights based on graph equivalence classes in [15] and dramatically reduces computation on real-world networks. The runtime vs. fidelity trade-off associated with this method is explored.

To obtain the inexact method, the spectral components are redefined as the generalized eigenspaces of adjacency matrix $A \in \mathbb{C}^{N \times N}$, and the resulting properties are discussed in Section II. Section III establishes the generalized Parseval's identity with respect to the inexact method. An equivalence class with respect to the inexact method is discussed in Section IV, which leads to a total variation ordering of the new spectral components as discussed in Section V. Section VI discusses the utility of the inexact method for real-world network analysis, and Section VII describes the trade-offs associated with its use.

Lastly, Section VIII demonstrates the application of the inexact GFT to New York City taxi data on the Manhattan street network. Our results demonstrate the speed of the inexact method and show that it reduces energy dispersal over spectral component corresponding to eigenvalue $\lambda = 0$.

II. Agile Inexact Method for Accelerating Graph Fourier Transform Computations

This section defines the Agile Inexact Method (AIM) for computing the graph Fourier transform (2) (GFT) of a signal. The *inexactness* of the AIM results from the abstraction of the Jordan subspaces that characterize (2); this concept is discussed further in this section. The *agility* of the AIM is related to the Jordan equivalence classes introduced in [15], which demonstrate that the degrees of freedom allowed by defective matrices yield GFT equivalence across networks of various topologies. Sections IV, VI, and VII discuss the increased computational efficiency that the AIM allows. We emphasize that reducing computation time is a major objective for formulating the inexact method.

The Agile Inexact Method (AIM) for computing the graph Fourier transform of a signal is defined as the signal projections onto the generalized eigenspaces of the adjacency matrix A. When A has k distinct eigenvalues, the definition of a generalized eigenspace $\mathscr{G}_i$ (see [13], [16]) is

$$\mathscr{G}_i = \mathrm{Ker}\,(A - \lambda_i I)^{m_i}\,,\ i = 1 \ldots, k \quad (3)$$

for corresponding eigenvalue λ_i and eigenvalue index m_i. Each $\mathscr{G}_i$ is the direct sum of the Jordan subspaces of λ_i.

The direct sum of generalized eigenspaces provides a unique decomposition of the signal space. This leads to the definition of the *Agile Inexact Method* for computing the graph Fourier transform of a graph signal $s \in \mathcal{S}$

$$\begin{aligned} \mathcal{H} : \mathcal{S} &\to \bigoplus_{i=1}^{k} \mathscr{G}_i \\ s &\to (\breve{s}_1, \ldots, \breve{s}_k)\,. \end{aligned} \quad (4)$$

The method when applied to signal s yields the unique decomposition

$$s = \sum_{i=1}^{k} \breve{s}_i, \qquad \breve{s}_i \in \mathscr{G}_i, \quad (5)$$

where each $\breve{s}_i$ is the (oblique) projection of s onto the ith generalized eigenspace $\mathscr{G}_i$ parallel to $\mathcal{S} \backslash \mathscr{G}_i$, or $\bigoplus_{j \neq i} \mathscr{G}_j$.

We define the projection operator for the inexact method. Denote eigenvector matrix

$$V = [v_{1,1} \cdots v_{1,a_1} \cdots v_{k,1} \cdots v_{k,a_k}], \quad (6)$$

where $v_{i,j}$ represents an eigenvector or generalized eigenvector of λ_i and a_i is the algebraic multiplicity of λ_i, $i = 1, \ldots, k$, $j = 1, \ldots, a_i$. Let V_i^0 represent the $N \times N$ matrix that consists of zero entries except for columns containing $v_{i,1}$ through v_{i,a_i}. The signal expansion components for signal s are $\widetilde{s} = Vs$ for signal $s \in \mathbb{C}^N$, where

$$\widetilde{s} = [\widetilde{s}_{1,1} \cdots \widetilde{s}_{1,a_1} \cdots \widetilde{s}_{k,1} \cdots \widetilde{s}_{k,a_k}]^T. \quad (7)$$

Therefore,

$$\begin{aligned} \breve{s}_i &= \widetilde{s}_{i,1} v_{i,1} + \cdots \widetilde{s}_{i,a_i} v_{i,a_i} && (8) \\ &= V_i^0 \widetilde{s} && (9) \\ &= V_i^0 V^{-1} s && (10) \end{aligned}$$

$$= V \begin{bmatrix} \ddots & & \\ & I_{a_i} & \\ & & \ddots \end{bmatrix} V^{-1} s \tag{11}$$

$$= Z_{i0} s, \tag{12}$$

where

$$Z_{i0} = V \begin{bmatrix} \ddots & & \\ & I_{a_i} & \\ & & \ddots \end{bmatrix} V^{-1} \tag{13}$$

is the *first component matrix* of the ith generalized eigenspace [13]. Missing entries in (11) and (13), including dot entries, are zero. The matrix Z_{i0} (13) projects $s \in \mathbb{C}^N$ onto the generalized eigenspace $\mathscr{G}_i$ parallel to $\bigoplus_{j \neq i} \mathscr{G}_j$. Note that the projection may be oblique. The reader is directed to Theorems 9.5.2 and 9.5.4 in [13] for more properties of (13).

The AIM provides a unique, Fourier transform-like projection space for a signal that is analogous to the definition (2) based on Jordan subspaces because the signal space is a direct sum of the generalized eigenspaces. The generalized eigenspaces are not necessarily irreducible components of the signal space, however, while the Jordan subspaces do provide an irreducible decomposition of the signal space. In terms of algebraic signal processing [17], [18], this means that the Jordan subspace formulation (2) can define a Fourier transform of a signal, while the generalized subspace (4) is inexact, and is not strictly a formulation of a true graph Fourier transform. The projections onto the generalized eigenspaces are additions of the spectral components of the GFT (2) and can be considered to be inexact analogues to the spectral components.

Note that the AIM (4) resolves to the Jordan subspace-based GFT (2) whenever the maximum number of Jordan subspaces per distinct eigenvalue is one (i.e., the characteristic and minimal polynomials of A are equal). Otherwise, Jordan subspace information is lost when the AIM is applied.

The following sections present key properties of (4), with particular emphasis on the generalized Parseval's identity, graph equivalence classes with respect to (4), and total variation ordering.

III. Generalized Parseval's Identity

Since the full graph Fourier basis V may not be unitary, the signal energy can be characterized by a generalized Parseval's identity. This identity is defined in [19] and used in [11] to characterize the spectral components. The identity is presented as a property below:

Property 1 (Generalized Parseval's Identity)**.** *Consider graph signal $s \in \mathbb{C}^N$ over graph $\mathcal{G}(A)$, $A \in \mathbb{C}^{N \times N}$. Let $V = [v_1 \cdots v_N]$ be a Jordan basis for A with dual basis matrix $W = V^{-H}$ partitioned as $[w_1 \cdots w_N]$. Let $s = \sum_{i=1}^{N} \langle s, v_i \rangle v_i = V \widetilde{s}_V$ be the representation of s in basis V and $s = \sum_{i=1}^{N} \langle s, w_i \rangle w_i = W \widetilde{s}_W$ be the representation of s in basis W. Then*

$$\|s\|^2 = \langle s, s \rangle = \langle \widetilde{s}_V, \widetilde{s}_W \rangle. \tag{14}$$

In this way, a biorthogonal basis set can be used to define signal energy. In particular, for a signal s over graph $\mathcal{G} = \mathcal{G}(A)$, V represents the eigenvector matrix of graph adjacency matrix A, and the columns of V and $W = V^{-H}$ form a dual basis. Thus, the expansions of signal s in these bases yield coefficient vectors $\widetilde{s}_V$ and $\widetilde{s}_W$ that satisfy (14).

We define the energy of a signal projection onto an AIM GFT component $\breve{s}_i$. Suppose $\breve{s}_i = \operatorname{span}(v_{i,1}, \ldots, v_{i,a_i})$, where $v_{i,1}, \ldots, v_{i,a_i}$ are columns of V with corresponding columns $w_{i,1}, \ldots, w_{i,a_i}$ of W. For signal s over graph $\mathcal{G} = \mathcal{G}(A)$, write $\breve{s}_i$ in terms of the columns of V as

$$\breve{s}_i = \alpha_1 v_{i,1} + \cdots + \alpha_{a_i} v_{i,a_i} \tag{15}$$

and in terms of the columns of W as

$$\breve{s}_i = \beta_1 w_{i,1} + \ldots \beta_{a_i} w_{i,a_i}. \tag{16}$$

Note that $\alpha = (\alpha_1, \ldots, \alpha_{a_i})$ and $\beta = (\beta_1, \ldots, \beta_{a_i})$ are subsets of coefficient vectors $\widetilde{s}_V$ and $\widetilde{s}_W$ that satisfy (14), respectively. The energy of $\breve{s}_i$ can then be defined as

$$\|\breve{s}_i\|^2 = \langle \alpha, \beta \rangle. \tag{17}$$

Equation (17) is used in Section VIII to compare the inexact method (4) and original GFT (2).

IV. Graph Equivalence under the Agile Inexact Method

Just as different choices of Jordan subspace bases yield the same GFT over Jordan equivalence classes of graph topologies in [11], different choices of bases for the generalized eigenspaces of the adjacency matrix yield the same inexact GFT. We define a $\mathscr{G}$-equivalence class as follows:

Definition 2 ($\mathscr{G}$-Equivalent Graphs)**.** *Consider graphs $\mathcal{G}(A)$ and $\mathcal{G}(B)$ with adjacency matrices $A, B \in \mathbb{C}^{N \times N}$. Then $\mathcal{G}(A)$ and $\mathcal{G}(B)$ are* $\mathscr{G}$-equivalent graphs *if all of the following are true:*

1) *The eigenvalues of A and B are identical (including algebraic and geometric multiplicities) – that is, $\mathcal{G}_A$ and $\mathcal{G}_B$ are cospectral; and*
2) *$\{\mathscr{G}_{A,i}\}_{i=1}^{k} = \{\mathscr{G}_{B,i}\}_{i=1}^{k}$, where $\mathscr{G}_{A,i}$ and $\mathscr{G}_{B_i}$ are the ith generalized eigenspaces of A and B, respectively, and k is the number of distinct eigenvalues.*

The set of all graphs that satisfy Definition 2 with respect to a graph adjacency matrix A with Jordan decomposition $A = VJV^{-1}$ form the $\mathscr{G}$-equivalence class of $\mathcal{G}(A)$. Denote by $\mathbf{G}_A$ the $\mathscr{G}$-equivalence class of $\mathcal{G}(A)$.

The power of the $\mathscr{G}$-equivalence classes comes from the observation that a graph adjacency matrix $A = VJV^{-1}$ may have the same eigenvalues and generalized eigenspaces as a matrix $B = \widetilde{V}J\widetilde{V}^{-1}$. For the AIM method (4), the basis provided by the columns of $\widetilde{V}$ yields identical signal projections as that provided by the columns of V. In addition, when $\widetilde{V}$ can be computed without finding Jordan chain vectors, the overall computation time for the eigendecomposition decreases. Section VII demonstrates the utility of these equivalence classes.

In contrast to the Jordan equivalence classes in [15], $\mathscr{G}$-equivalence classes are independent of Jordan subspaces and their dimensions. This allows for greater degrees of freedom in basis choices that yield the same signal projections. As a result, the eigendecomposition required for the AIM (4) can be simpler and faster to compute than that for the GFT (2). A particular case is discussed in greater detail in Section VII.

The next section defines a total variation ordering over the generalized eigenspaces $\mathscr{G}_i$ (3) of the graph adjacency matrix. A class total variation is proposed to account for different choices of Fourier bases that correspond to a particular $\mathscr{G}$-equivalence class.

V. Total Variation Ordering

In [11] and [15], we defined an ordering on the GFT spectral components based on the total variation of a signal. This allows ordering the frequency components so as to define low-pass, band-pass, and high-pass graph signals with respect to specific graphs. Using a total variation-based characterization signifies that the variation between a signal and its transformation by the graph adjacency matrix is less for low-pass signals than for high-pass signals. In this section, such an ordering is shown for GFT spectral components corresponding to the Agile Inexact Method (AIM).

Denote by V_i the submatrix of columns of eigenvector matrix V that span $\mathscr{G}_i$, the ith generalized eigenspace. The (graph) total variation of V_i is defined as

$$\mathrm{TV}_G\left(V_i\right) = \left\|V_i - AV_i\right\|_1. \tag{18}$$

The definition (18) of total variation has been characterized in [11] where each V_i spans a Jordan subspace instead of a generalized eigenspace. Since the generalized eigenspaces are direct sums of Jordan subspaces (see Section I), many properties hold for the total variation over generalized eigenspaces. Theorem 3 shows that the total variation is invariant to a permutation of node labels. Theorem 4 shows an upper bound on this total variation. These theorems follow immediately from the proofs of Theorems 25 and 28 of [11]. These theorems show the dependence of the total variation (18) on the choice of eigenvector matrix V and motivate the definition of basis-independent class total variations as in [15].

Theorem 3. *Let $A, B \in \mathbb{C}^{N\times N}$ such that $\mathcal{G}(B)$ is isomorphic to $\mathcal{G}(A)$. Let $V_{A,i} \in \mathbb{C}^{N\times a_i}$ be a union of Jordan chains that span generalized eigenspace $\mathscr{G}_i$ of matrix A and $V_{B,i} \in \mathbb{C}^{N\times a_i}$ the corresponding union of Jordan chains of B. Then $\mathrm{TV}_G(V_{A,i}) = \mathrm{TV}_G(V_{B,i})$.*

We briefly note that the generalized eigenspaces of the adjacency matrices A and B in Theorem 3 are not necessarily equivalent but are related by an isomorphism. The reader is referred to [15] for more on isomorphic equivalence classes.

Theorem 4. *Consider adjacency matrix A with k distinct eigenvalues and $N \times a_i$ eigenvector submatrices V_i with columns corresponding to the union of the Jordan chains of λ_i, $i = 1, \ldots, k$. Then the graph total variation is bounded as*

$$\mathrm{TV}_G(V_i) \leq |1 - \lambda_i| + 1. \tag{19}$$

The bound (19) highlights that the choice of Jordan chain vectors in the eigendecomposition of the adjacency matrix affects the total variation. In order to gain independence from a choice of basis, the definition of total variation can be generalized to a class total variation defined over the $\mathscr{G}$-equivalence class associated with the graph of adjacency matrix A. We define the *class total variation* of spectral component $\mathscr{G}_i$ as the supremum of the graph total variation of V_i over the $\mathscr{G}$-equivalence class (for all $\mathcal{G}(B) \in \mathbf{G}_A$):

$$\mathrm{TV}_{\mathbf{G}_A}\left(\mathscr{G}_i\right) = \sup_{\substack{\mathcal{G}(B)\in\mathbf{G}_A \\ B=VJV^{-1} \\ \mathrm{span}\{V_i\}=\mathscr{G}_i \\ \|V_i\|_1=1}} \mathrm{TV}_G\left(V_i\right). \tag{20}$$

In particular, the low-to-high variation sorting is achieved via the function $f(\lambda_i) = |1 - \lambda_i| + 1$.

VI. Applicability to Real-World Networks

This section discusses how the AIM can be applied on real-world large, sparse, and directed networks. We discuss examples that illustrate that adjacency matrices for these networks have a single generalized eigenspace of dimension greater than one. We demonstrate the ease of computing the AIM for such networks.

A. Large, Directed, and Sparse Networks

In practice, sparsity in directed networks yields a zero eigenvalue of high algebraic and geometric multiplicities

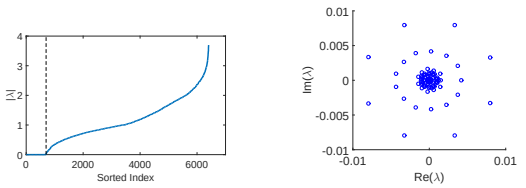

(a) Eigenvalue magnitudes for road network

(b) Small eigenvalues for road network

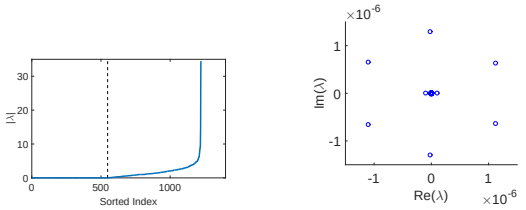

(c) Eigenvalue magnitudes for political blog

(d) Small eigenvalues for political blog

Fig. 1: (a) Eigenvalue magnitudes of a directed New York City road network adjacency matrix. The magnitudes are small in the index range 1 to 699 (up to the dotted line). (b) Road network eigenvalues (699 total) in the sorted index range 1 to 699 plotted on the complex plane. (c) Eigenvalue magnitudes of the largest weakly connected component of a political blog network. The magnitudes are small in the index range 1 to 548 (up to the dotted line). (d) Blog eigenvalues (548 total) in the sorted index range 1 to 548.

that may not be equal. This behavior can be illustrated using two examples. The first is a directed, strongly connected road network $\mathcal{G}_{\rm rd} = \mathcal{G}(A_{\rm rd}) = (\mathcal{V}_{\rm rd}, \mathcal{E}_{\rm rd})$ of Manhattan, New York [20]; $\mathcal{G}_{\rm rd}$ contains 6,408 nodes that represent the latitude and longitude coordinates of intersections and mid-intersection locations (available from [21]), as well as 14,418 directed edges, where edge $e = (v_1, v_2)$ represents a road segment along which traffic is legally allowed to move from v_1 to v_2 as determined from Google Maps [20]. The second network example is the largest weakly connected component $\mathcal{G}_{\rm bl} = \mathcal{G}(A_{\rm bl}) = (\mathcal{V}_{\rm bl}, \mathcal{E}_{\rm bl})$ of a political blog network, with 1,222 blogs as nodes and 19,024 directed edges, where edge $e = (v_1, v_2)$ exists between blogs $v_1, v_2 \in \mathcal{V}_{\rm bl}$ if v_1 contains a link to v_2 [22].

Figures 1b and 1d show that the numerical eigenvalues of these real-world networks occur in clusters on the complex plane. These eigenvalues have small magnitude (lying left of the dashed lines in Figures 1a and 1c). On first inspection, it is not obvious whether the numerical eigenvalues are spurious eigenvalues clustered around a true eigenvalue of zero, or whether the numerical eigenvalues are indeed the true eigenvalues.[1] This phenomenon is typically observed in systems with multiple eigenvalues [23], [24], [25]. In practice, we can verify that numerical zero is an eigenvalue using either pseudospectra [26] or techniques that explore the numerical stability of the singular value decomposition [27], [28]. For the examples explored here, the singular value decomposition demonstrated that the clusters of low-magnitude eigenvalues represented a numerical zero eigenvalue. This method involves comparing small singular values to machine precision and is discussed in more detail in Section VIII-C.

Furthermore, the adjacency matrices $A_{\rm rd}$ and $A_{\rm bl}$ have kernels of dimension 446 and 439, respectively[2], confirming that eigenvalue $\lambda = 0$ has high algebraic (and geometric) multiplicity. In addition, eigenvector matrices $V_{\rm rd}$ and $V_{\rm bl}$ computed with MATLAB's eigensolver have numerical rank $6363 < 6408$ and $910 < 1222$, respectively, implying the existence of nontrivial Jordan subspaces (Jordan chains of length greater than one) for $\lambda = 0$. While these general eigenspace properties can be readily determined, a substantial amount of additional computation is required to determine the dimensions of each Jordan subspace for eigenvalue $\lambda = 0$. For example, the Jordan subspaces for $\lambda = 0$ can be deduced by computing a generalized null space as in [30], but this computation takes $O(N^3)$ or $O(N^4)$ for an $N \times N$ matrix.

On the other hand, the eigenvectors corresponding to eigenvalues of higher magnitude in these networks are full rank. In other words, from a numerical perspective, the Jordan subspaces for these eigenvalues are all one-dimensional.

In such networks, for which sparsity yields a high-multiplicity zero eigenvalue while the other eigenvalues correspond to a full-rank eigenvector submatrix, the AIM inexact method given in (4) can be applied in the following way. Denote the eigenvector matrix V of graph adjacency matrix A such that

$$V = \begin{bmatrix} V_{\rm known} & \widetilde{V} \end{bmatrix}, \tag{21}$$

where $V_{\rm known}$ be the full-rank eigenvector submatrix corresponding to the nonzero eigenvalues, and the columns of submatrix $\widetilde{V}$ span the generalized eigenspace $\mathcal{G}_0$ corresponding to eigenvalue zero.

One way to find a spanning set of columns for $\widetilde{V}$ is to find the kernel of $V_{\rm known}^T$, since the range space of The

[1] The eigenvalues in Figure 1 were computed in MATLAB. Similar clusters appear using `eig` without balancing and `schur`. They also appear when analyzing row-stochastic versions of the asymmetric adjacency matrices.

[2] The kernels (null spaces) were computed in MATLAB [29] on a 16-GB, 16-core Linux machine.

resulting inexact graph Fourier basis is then

$$\widehat{V} = \begin{bmatrix} V_{\text{known}} & \operatorname{Ker}\left(V_{\text{known}}^T\right) \end{bmatrix}. \tag{22}$$

If the original network has adjacency matrix $A = VJV^{-1}$, then (22) implies that the network with adjacency matrix $\widehat{A} = \widehat{V}J\widehat{V}^{-1}$ is in the same $\mathscr{G}$-equivalence class as the original network with adjacency matrix $A = VJV^{-1}$. As discussed in Section IV, the AIM GFT is equivalent over both networks. Section VII shows that computing the Fourier basis as in (22) is computationally fast and efficient.

In this way, for large, sparse, and directed real-world networks that exhibit a single zero eigenvalue of high multiplicity, the AIM can simplify the computation of a graph Fourier basis. Instead of computing Jordan chains or generalized null spaces to deduce the structure of the Jordan subspaces, a single kernel computation is needed to compute the spanning basis. This idea is explored further in Section VII.

B. Multiple Nontrivial Jordan Subspaces

In the case of multiple distinct eigenvalues with large but unequal algebraic and geometric multiplicities, it may be possible to compute a few Jordan chain vectors to differentiate the nontrivial generalized eigenspaces. This approach is only computationally efficient if the number of Jordan chain vectors to compute is relatively small, as discussed in [14], [31].

When this is not feasible, a coarser inexact method can be obtained by reformulating (22) so that V_{known} consists of all known eigenvectors and is full rank. Then a basis of the kernel of V_{known}^T can be computed to obtain a set of spanning vectors.

It is unknown a priori which basis vectors of $\operatorname{Ker} V_{\text{known}}^T$ belong to the nontrivial generalized eigenspaces of A. This can be overcome by randomly assigning these vectors to the nontrivial generalized eigenspaces, which results in a coarser inexact method. The dimensions of each generalized eigenspace must be known so that the correct number of vectors is assigned to each eigenspace. The dimensions can be determined by recursively computing for eigenvalue λ_i

$$f(l) = \dim \operatorname{Ker}(A-\lambda_i I)^l - \dim \operatorname{Ker}(A-\lambda_i I)^{l-1} \tag{23}$$

for $l = 2, \ldots, N$; this equation provides the number of Jordan chains of length at least l [13]. The value of $\dim \operatorname{Ker}(A-\lambda_i I)^{l-1}$ for $l > 1$ at which $f(l) = 0$ is the dimension of generalized eigenspace $\mathscr{G}_i$. If $|\lambda_{\max}| > 1$, the condition $f(l) = 0$ may not be attained; instead, $f(l)$ becomes a monotonically increasing function for large l. In this case, the value of $\dim \operatorname{Ker}(A - \lambda_i I)^{l-1}$ for l at which $f(l)$ becomes monotonically increasing is the approximate generalized eigenspace dimension.

The random assignment of missing eigenvectors to the nontrivial generalized eigenspaces would be a coarser inexact method compared to the AIM method (4); in particular, it is unknown which eigenvector assignment, if any, corresponds to a graph in the same $\mathscr{G}$-equivalence class as the original graph. On the other hand, the matrices built from a series of random assignments could be used to construct a filter bank such that each assignment corresponds to a different graph filter. An optimal or near-optimal weighted combination of such filters could be learned by using time-series of filtered signals as inputs to train a classifier. While the resulting Fourier basis is only an approximation of the one required for the AIM method (4), such a learned combination of filters would provide an interpretable tool to analyze graph signals while also ranking each filter (representing a random assignment of vectors to generalized eigenspaces) in terms of its utility for typical graph signals.

The next section characterizes the loss of fidelity to the original graph when the AIM is applied as in (22) for a single unknown, nontrivial generalized eigenspace.

VII. Runtime vs. Fidelity Trade-off

While maximizing the fidelity of the network representation is desirable in general, obtaining as much knowledge of the Jordan chains as possible is computationally expensive. As shown below, execution time is linear with respect to the number of Jordan chain vectors to compute. In practice, however, computation of increasing numbers of chains requires allocating memory for matrices of increasing size. This memory allocation time can increase nonlinearly once specific limits (e.g., RAM) imposed by the computing hardware are exceeded. Since practical applications of GFT analysis may be constrained by execution time requirements and available computing hardware, methods that allow fidelity vs. speedup tradeoffs will prove useful.

Our inexact GFT approach provides precisely such a method for trading between higher fidelity analysis (a greater number of Jordan chains computed) and execution time, as discussed in Section VI. In this section, the details of such trade-offs are discussed in detail.

To illustrate the cost of computing the Jordan chain vectors of $A \in \mathbb{C}^{N\times N}$ that complete the generalized eigenspace of eigenvalue λ_i, consider the general algorithm given by:

```
1: function COMPUTE_CHAINS
2:     V = FULL_RANK_EIGENVECTORS(A)
3:     R = PSEUDOINVERSE(A − λ_i I)
4:     N = KER(A − λ_i I)
5:     Nnew=[ ]
6:     for c in 1..maximum chain length do
7:         for v in N do
8:             if [(A − λ_i I)  v] full rank then
```

```
            v2 = R*v
            Append v2 to V.
            Append v2 to Nnew.
          end if
        end for
        N = Nnew
      end for
    end function
```

The eigenvector and pseudoinverse computations are pre-processing steps that can be executed just once for a network with stationary topology. Efficient methods for these computations include power iterations, QR algorithms, and, for large, sparse matrices, Lanczos algorithms and Krylov subspace methods [14]. The key bottlenecks of concern reside in the `for` loop. This loop consists of a rank-checking step to verify that the current eigenvector is in the range space of $A-\lambda_i I$. This step can be implemented using the singular value decomposition, which has $O(kMN^2 + k'M^3)$ operations [14] on an $M \times N$ matrix, $M > N$; constants k and k' could be 4 and 22, respectively, as in the R-SVD algorithm of [14]. In addition, the matrix-vector product in the `for` loop takes about $2MN$ operations [14]. The original dimension of V is $M \times N(j)$, where the number of columns $N(j)$ approaches M from below as the number j of vectors traversed increases. The resulting time complexity for a single iteration is

$$O(kMN(j)^2 + k'M^3) + O(2MN(j)). \quad (24)$$

The first `for` loop in `compute_chains` iterates m_i times, where m_i is the maximum chain length corresponding to λ_i. A method for estimating m_i is demonstrated in Section VIII-D. The second `for` loop will first run for g_i iterations, where g_i is the kernel dimension of $A - \lambda_i I$ (the geometric multiplicity of λ_i). On the lth run of the outer loop, the number of iterations of the inner loop equals the difference of kernel dimensions $f(l)$ given by (23). Therefore, the total number b_A of iterations is $m_i \sum_{l=1}^{m_i} f(l)$, or

$$m_i(\dim \mathrm{Ker}(A - \lambda_i I)^{m_i} - \dim \mathrm{Ker}(A - \lambda_i I)) \quad (25)$$
$$= m_i(a_i - g_i), \quad (26)$$

where a_i and g_i are the algebraic and geometric multiplicities of λ_i, respectively. Since b_A (26) depends on the adjacency matrix A, it has an implicit dependence on the $M \times N(j)$ dimensions of A. The total time complexity of the `for` loops is then

$$\sum_{j=1}^{b_A} O(kMN(j)^2 + k'M^3) + O(2MN(j)). \quad (27)$$

In addition, the time to allocate memory for the matrix $[(A - \lambda I)\ \ \mathrm{v}]$ is

$$O(MN(j)) \cdot m(MN(j)), \quad (28)$$

where $m(\cdot)$ is the platform-dependent time per unit of memory allocation as a function of the matrix size. Since each `for` loop allocates memory in this way, the total memory allocation time is

$$\sum_{j=1}^{b_A} O(MN(j)) \cdot m(MN(j)). \quad (29)$$

Assuming c is the platform-dependent time per floating-point operation and SVD constants $k = 4$ and $k' = 22$, the total expected runtime for the Jordan chain computation can be approximated as

$$\sum_{j=1}^{b_A} c(4MN(j)^2 + 22M^3 + 2MN(j)) + MN(j)m\left(MN(j)\right), \quad (30)$$

where c is the platform-dependent time per floating-point operation and the full SVD complexity coefficients are set to $k = 4$ and $k' = 22$.

Equation (30) shows that the runtime for the Jordan chain computations is approximately linear with the number of missing vectors one needs to compute. In practice, however, the time to allocate memory can scale nonlinearly with the number of nodes as the size of the eigenvector matrix approaches system capacity. For networks at massive scale, this can present a significant problem.

In contrast, the algorithm to compute an orthogonal set of missing vectors is very fast. The algorithm is the following:

```
function COMPUTE_MISSING
    V = FULL_RANK_EIGENVECTORS(A)
    Vt = TRANSPOSE(A)
    Vnew = KER(Vt)
    V = [V  Vnew]
end function
```

The bottleneck in `compute_missing` is computing the kernel of V^T, which has time complexity $O(kMN(j)^2 + k'M^3)$ for a SVD-based implementation. In addition, memory allocation for Vt and V each takes $O(MN(j))$ time. Notably, this allocation is only performed once in the AIM implementation (22) ($O(1)$ in space), as compared to the m_i allocations required in the `for` loop of `compute_chains`.

While the output V matrix in `compute_missing` yields the same AIM results given by (4) as the original eigenvector matrix, the corresponding adjacency matrix $\widetilde{A}$ may not preserve key structures in the original graph adjacency matrix A. In other words, the trade-off

for the faster computation time in `compute_missing` is a loss of fidelity to the original graph. In order to improve the fidelity, one may compute a certain number $b < b_A$ of Jordan chains as determined by RAM and other system constraints. In this way, the AIM can be used to selectively trade off execution time (30) against fidelity in the computations.

This section shows that the AIM can be employed to enable GFT analysis to meet execution time constraints by trading runtime for fidelity. The analysis of algorithms `compute_chains` and `compute_missing` quantifies the expected execution time in terms of a given number of Jordan chain vectors. These considerations need to be addressed when applying the graph Fourier transform (2) and the AIM method (4) to real-world problems. Such an application is shown in the next section.

VIII. New York City Taxi Data: Eigendecomposition and Graph Fourier Transform

This section expands on the issues described in Section VI and presents details for computing the eigenvector matrix V for the non-diagonalizable adjacency matrix A of the Manhattan road network.

The Agile Inexact Method (AIM) for the graph Fourier transform developed in Section II is then applied to the statistics computed from four years of New York City taxi data. Our method provides a fine-grained analysis of city traffic that should be useful for urban planners, such as in improving emergency vehicle routing.

Sections VIII-B– VIII-E describe our method to find the eigendecomposition of the Manhattan road network. Section VIII-F demonstrates the AIM for computing the graph Fourier transform.

A. Data Set Descriptions

Our goal is to extract behaviors over space and time that characterize taxi movement through New York City based on four years (2010-2013) of New York City taxi data [32]. Since the path of each taxi trip is unknown, an additional processing step is required to estimate taxi trajectories before extracting statistics of interest. For example, if a taxi picks up passengers at Times Square and drops them off at the Rockefeller Center, it is desirable to have statistics that capture not just trip data at the landmarks, but also at the intermediate locations.

Estimating tax trajectories requires overlaying the taxi data on the New York City road network. Details on the taxi data and the road network are described as follows.

1) NYC Taxi Data: The 2010-2013 taxi data we work with consists of 700 million trips for 13,000 medallions [32]. Each trip has about 20 descriptors including pick up and drop off timestamps, latitude, and longitude, as well as the passenger count, trip duration, trip distance, fare, tip, and tax paid. The data is available as 16.7 GB of compressed CSV files.

Dynamic representation. Since the available data provides only static (start and end) information for each trip, an estimate of the taxi trajectory is needed in order to capture the taxi locations interspersed throughout the city at a given time slice. Our method for estimating these trajectories, which is based on Dijkstra's algorithm, is described in detail in [33].

Once taxi paths are estimated, they are used to extract traffic behavior. The statistic of interest we consider here is the average number of trips that pass through a given location at a given time of day.

2) NYC and Manhattan Road Networks: The road network $G = (V, E)$ consists of a set V of $|V| = 79,234$ nodes and a set E of $|E| = 223,966$ edges which we represent as a $|V| \times |V|$ adjacency matrix A. The nodes in V represent intersections and points along a road based on geo-data from [21]. Each edge $(v_i, v_j) \in E$, $v_i, v_j \in V$, corresponds to a road segment on which traffic may flow from geo-location v_i to geo-location v_j as determined by Google Maps [34]. An edge of length $d_{ij} > 0$ is represented by a nonzero entry in the adjacency matrix so that $[A]_{ij} = d_{ij}$. A zero entry corresponds to the absence of an edge.

The network G is *directed* since the allowed traffic directions along a road segment may be asymmetric. In addition, G is *strongly connected* since a path (trajectory) exists from any geo-location to any other geo-location.

Our analysis focuses on the Manhattan grid, which is a subgraph of G with $6,408$ nodes and $14,418$ edges. This subgraph is also strongly connected. The eigendecomposition in Section VIII-B is based on this network. Expressing the average number of trips that pass through a given location at a fixed time of day as a vector over the Manhattan grid defines a graph signal to which the GFT (2) and AIM method (4) can be applied.

B. Eigendecomposition of Manhattan Road Network

The Manhattan road network described in Section VIII-A provides the adjacency matrix $A \in \mathbb{R}^{6408\times 6408}$ for which we compute the eigendecomposition. For this particular adjacency matrix, the eigenvector matrix $V_{\rm obs} \in \mathbb{C}^{6408\times 6408}$ generated by a standard eigenvalue solver such as MATLAB or LAPACK is not full rank, where the rank can be computed with either the singular value decomposition or the QR decomposition; i.e., A is not diagonalizable.

In order to apply the GFT (2) or the original projections onto each eigenvector as in [1], Jordan chain computations are required. These computations are costly as described in Section VII and also potentially numerically unstable. In contrast, the AIM (4) does not require this step.

Since it is instructive to compare the GFT with the AIM results, it is necessary to compute the Jordan chains here. For this reason, the following sections detail our method for computing the Jordan chains for the Manhattan road network.

C. Eigenvalue Verification and Maximum Jordan Chain Length

Since the kernel dimension of A (the geometric multiplicity of the zero eigenvalue) is 446, and there are 699 eigenvalues of magnitude close to zero (as seen in Figure 1a), $699 - 446 = 253$ generalized eigenvectors must be determined to complete the Fourier basis.

As discussed in Section VI-A, there is ambiguity in determining whether the 699 eigenvalues of small magnitude represent true zero eigenvalues. Consequently, accurately identifying eigenvalues of high algebraic multiplicity can be a hard problem – see, for example, [24], [25], [26], [27], [31], [14]. Furthermore, these eigenvalues of small magnitude form constellations with centers close to zero, as shown in Figure 1b. References [24] and [25] show that the presence of such constellations may indicate the presence of an eigenvalue of high algebraic multiplicity; [24] states that an approximation of this eigenvalue may be the center of the cluster of numerical eigenvalues, while [25] presents an iterative algorithm to approximate a "pseudoeigenvalue." It is unclear from pure inspection[3] whether the observed constellations are a result of the multiplicity of a zero eigenvalue or are the actual eigenvalues of A. For these reasons, it is necessary to verify the existence of a numerical zero eigenvalue before computing the Jordan chains. Our method is explained below; see also [20].

To verify a numerical zero eigenvalue, a result from [27] is applied, which states the following:

Result 5 ([27])**.** *Consider a matrix $A \in \mathbb{C}^{N\times N}$ with singular values $\sigma_1(A) > \cdots > \sigma_N(A)$ that is scaled so that $\sigma_1(A) = 1$. Let $m_k = \left|\mathrm{Ker}(A^k)\right|$ denote the dimension of the kernel of A^k. In addition, let α and δ be positive constants; δ is usually on the order of machine precision and α is significantly greater than δ. Then 0 is a numerically multiple eigenvalue with respect to α and δ if*

$$\sigma_{N-m_k}\left(A^k\right) > \alpha > \delta > \sigma_{N-m_k+1}\left(A^k\right), \qquad (31)$$

for $k = 1, 2, \ldots, h$, where h is the maximum Jordan chain length for the zero eigenvalue.

Since the constants α and δ have different orders of magnitude, Equation (31) implies that singular value $\sigma_{N-m_k}(A^k)$ is significantly greater than $\sigma_{N-m_k+1}(A^k)$.

Result 5 serves two purposes for our application. First, it verifies the existence of a numerical zero eigenvalue. It also implies that a value of k at which Equation (31) fails cannot be the maximum Jordan chain length of the zero eigenvalue. Accordingly, we propose the following method to find the maximum Jordan chain length h: increment the value of k starting from $k = 1$, and let $k = k'$ be the first value of k such that Equation (31) fails. Then the maximum Jordan chain length for eigenvalue zero is $h = k' - 1$.

Table I is obtained by applying Result 5 to the adjacency matrix A of the Manhattan road network. The columns of the table correspond to the power k of A, the dimension m_k of the kernel of A^k, the index $N - m_k$ of the first singular value of A^k to examine, and the values of the singular values at indices $N-m_k$ and $N-m_k+1$. The results in the table are reasonable since the computational machine precision was on the order of 10^{-16}, and the first four rows of the table display singular values for which δ is on the order of 10^{-15} or 10^{-16} and the constant α is larger by 6 to 12 orders of magnitude. Therefore, the inequality (31) holds for the first four rows of Table I. The inequality begins to fail at $k = 5$; thus, the expected maximum numerical Jordan chain length is no more than 3 or 4.

D. Jordan Chain Computation

The steps for finding the eigenvectors for eigenvalue zero are described here. This is done to provide ground truth to assess the application of the AIM (4) to real data as shown in Section VIII-F.

To find the eigenvectors for eigenvalue zero, the null space or kernel $\mathrm{Ker}(A)$ of A is first computed to find the corresponding eigenvectors. Each of these eigenvectors corresponds to a Jordan block in the Jordan decomposition of A, and the Jordan chains with maximum length h (as determined by applying Result 5 from [27] as shown in Table I) are computed by the recurrence equation [13]

$$Av_k = \lambda v_k + v_{k-1} = v_{k-1}, \qquad (32)$$

where $k = 1, \ldots, h$ and $v_0 \in \mathrm{Ker}(A)$. If the number of linearly independent proper and generalized eigenvectors

[3]Considering double-precision floating point (64 bit) and that the number of operations to compute the eigenvalues of an $N \times N$ matrix is $O(N^3)$, the expected precision is on the order of 10^{-6} or 10^{-7}. Numerous eigenvalues in Figures 1a and 1b demonstrate this order of magnitude.

k	m_k	$N - m_k$	$\sigma_{N-m_k}\left(A^k\right)$	$\sigma_{N-m_k+1}\left(A^k\right)$
1	446	5962	1.9270×10^{-3}	1.2336×10^{-15}
2	596	5812	2.1765×10^{-6}	6.9633×10^{-16}
3	654	5754	1.4013×10^{-8}	3.4250×10^{-16}
4	678	5730	1.1853×10^{-10}	3.1801×10^{-16}
5	692	5716	2.0163×10^{-11}	8.4063×10^{-14}
6	700	5708	9.6533×10^{-11}	8.2681×10^{-11}

TABLE I: Singular values to validate existence of a numerical zero.

equals N, we are done; otherwise, the Fourier basis must be extended.

Numerical instability. Numerically stable methods such as the power method, QR algorithms, and Jacobi methods for diagonalizable matrices are applicable methods for eigendecomposition; see [14] and the references therein for more details. Furthermore, eigendecompositions of large, sparse matrices can be computed with iterative Lanczos or Krylov subspace methods [14, Chapter 10], [35].

On the other hand, for defective or nearly defective matrices, small perturbations in network structure due to numerical round-off errors can significantly perturb the numerical eigenvalues and eigenvectors; for example, for a defective eigenvalue of A corresponding to a p-dimensional Jordan block, $O(\epsilon)$ perturbations in A can result in $O(\epsilon^{1/p})$ perturbations in λ [14], [36]. In addition, while computing the Jordan normal form can be numerically unstable [27], [28], [31], [14], forward stability was shown for the Jordan decomposition algorithm in [37] and forms the basis for an SVD-based implementation here. The Jordan chains are generated from the vectors in the kernel of A^h, where h is the maximum Jordan chain length; these computations are based on the SVD or the QR decomposition, and so are more stable. Then the Jordan chain can be constructed by direct application of the recurrence equation (32).

Computation details. A single pass of the Jordan chain algorithm requires about a week to run on a 30-machine cluster of 16 GB RAM/16-core and 8 GB RAM/8-core machines; however, the algorithm does not return a complete set of 253 generalized eigenvectors. To maximize the number of recovered generalized eigenvectors, successive computations were run with different combinations of vectors that span the kernel as starting points for the Jordan chains. After three months of testing, the best run yielded 250 of the 253 missing Jordan chain vectors; the remaining eigenvectors were computed as in (22). The maximum Jordan chain length is two, which is consistent with Result 5 and Table I. As a result, the constructed eigenvector matrix captures a large part of the adjacency matrix structure. The next section describes our inexact alternative to full computation of the Jordan chains, which provides a significant speedup in the determination of a useful basis set.

E. Generalized Eigenspace Computation

In order to compute the generalized eigenspace for the zero eigenvalue, the known eigenvectors are determined as in Section VIII-D. Then the eigenvector matrix is computed as (22). This is a five minute computation on a 64-bit machine with 16 cores and 16 GB RAM.

F. Demonstration of the Agile Inexact Method

In this section, the original graph Fourier transform, defined as eigenvector projections, is compared to the Agile Inexact Method (AIM) (22) for a graph signal representing the four-year average of trips that pass through Manhattan from June through August on Fridays from 9pm to 10pm.

Figures 2a and 2b show the percentage of energy that is contained in the signal projections onto eigenvectors and onto generalized eigenspaces, respectively, using our previously defined energy metric (Section III). The maximum energy contained in any of the eigenvector projections is 1.6% as shown in the data point in Figure 2a. The observed energy dispersal in the eigenvector projections corresponding to $\lambda = 0$ is a result of oblique projections. The computed Jordan chain vectors are nearly parallel to the proper eigenvectors (angles less than 1 degree), so the signal projection onto a proper or generalized eigenvector v of $\lambda = 0$ parallel to $\mathbb{C}^N \backslash \mathrm{span}(v)$ has an augmented magnitude compared to the orthogonal projection. The reader is referred to [38] for more on oblique projectors.

In contrast, projecting onto the generalized eigenspaces with the AIM method (4) concentrates energy into two eigenspaces. The data points in Figure 2b at indices 706 and 707 contain 30% and 28% of the signal energy, respectively. On the other hand, the energy concentrated on the generalized eigenspace for

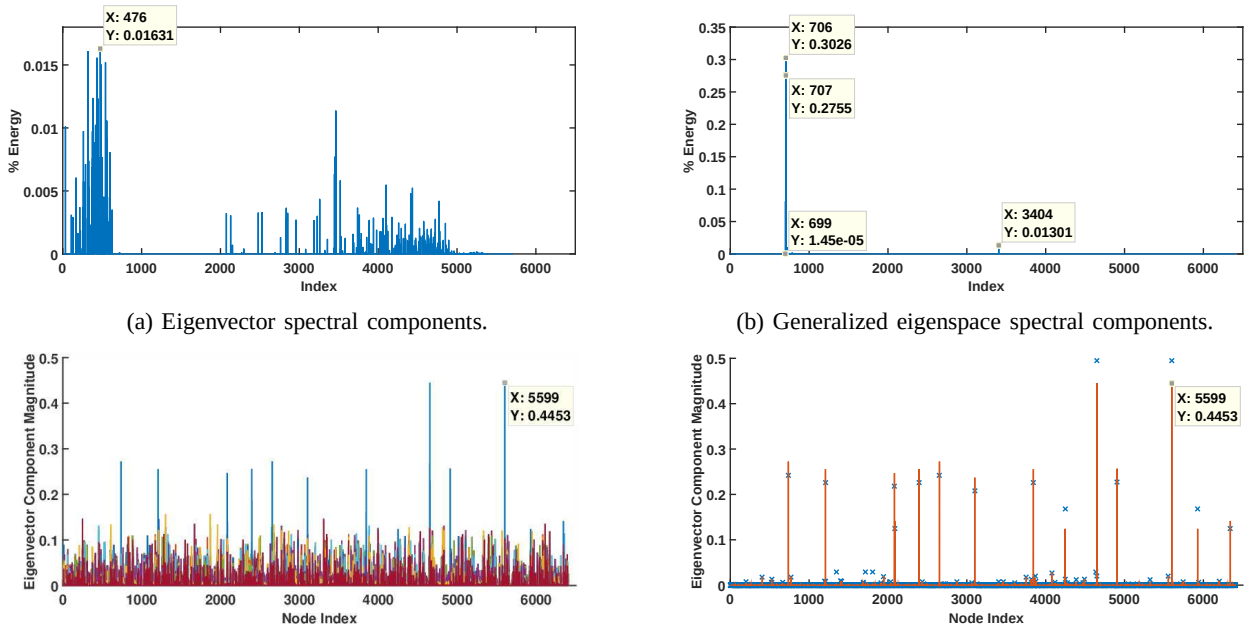

(a) Eigenvector spectral components.

(b) Generalized eigenspace spectral components.

(c) Eigenvector component magnitudes.

(d) Eigenvector component magnitudes.

Fig. 2: (a) Percentage of energy captured by signal projections onto eigenvectors, including generalized eigenvectors. The horizontal axis corresponds to the eigenvector of the Manhattan road matrix based on the ascending order of the corresponding eigenvalue magnitudes. (b) Percentage of energy captured by signal projections onto generalized eigenspaces. The horizontal axis corresponds to the generalized eigenspace of the Manhattan road matrix based on the ascending order of corresponding eigenvalue magnitudes. The first projection energy onto the generalized eigenspace of $\lambda = 0$ is shown at index 699, corresponding to less than 1% of the total signal energy. The preceding indices are set to zero to match the indices of (a). The following indices correspond to the same eigenvectors as in (a). The most energy is captured by projections onto the eigenvector at index 706 (30% of the energy) and at index 707 (28% of the energy). (c) Eigenvectors of (a) that capture 60% of the signal energy. The horizontal axis corresponds to a node (geo-coordinate) in the Manhattan road network. The vertical axis corresponds to the magnitude of the corresponding eigenvector component. There are 84 eigenvectors corresponding to about 60% of the signal energy, each shown in a different color. (d) The two eigenvectors in (b) that correspond to 58% of the total signal energy, one depicted as red lines and the other as blue x's. The axes are identical to those in (c). The eigenvector components of highest magnitude correspond to the same nodes, or locations in New York City. Furthermore, the locations with highest eigenvector expression match those in (c).

eigenvalue zero is less than 1%, as shown by the data point at index 699 in Figure 2b. Thus, the AIM method decreases energy dispersal over the spectral components and reduces the effect of oblique projections shown in Figure 2a.

For both methods (the eigenvector projections and the AIM), the spectral components are arranged by the energies they contain, in decreasing order, and those that contain 60% of the signal energy in the projections have their component magnitudes plotted in Figures 2c and 2d. There are 84 such eigenvectors for the projections onto eigenvectors, compared to only two eigenspaces (eigenvectors in this case) that contain the most energy in the generalized eigenspace case; this illustrates that the eigenvector projections disperse more energy than the generalized eigenspace projections. Figures 2c and 2d show that the highly expressed eigenvector components are located at the same nodes for both methods. This demonstrates that the AIM can provide similar results to the original formulation, with a significant acceleration of execution time and less energy dispersal over the spectral components.

The highly expressed components of eigenvectors 706 and 707 in Figure 2d are plotted on the Manhattan grid in Figure 3b. There are two clusters on the east side of Manhattan, one in black in Gramercy Park and one in purple by Lenox Hill. These locations represent sites of significant behavior that suggest concentrations of taxi activity. The corresponding graph signal, shown in Figure 3a, confirms this. Since this signal represents the average number of taxi trips on Fridays from 9pm to 10pm, one possible explanation for this behavior is the

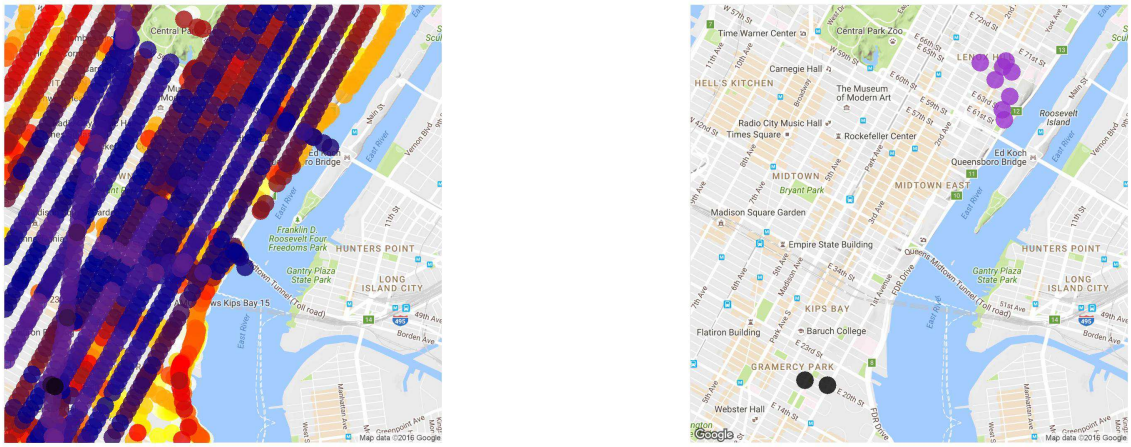

(a) Original signal.

(b) Maximum eigenvector components.

Fig. 3: (a) Original signal, representing the June-August four-year average number of trips through each node on the Manhattan road network. Colors denote 699 log bins (white: 0–20, beige to yellow: 20–230, orange: 230–400, red: 400–550, blue: 550–610, purple to black: 610–2,700). (b) Highly expressed eigenvector in Figures 2c and 2d, where the maximum components and second maximum components are represented by black and purple dots, respectively. These clusters correspond to locations through which high volumes of taxi trips flow.

proximity of the expressed locations to restaurants and eateries that are known gathering spots for taxi drivers. Visual comparison of Figure 3a and 3b highlights the utility of our method in finding fine-grained behaviors that are otherwise non-obvious from the raw signal.

This section provides a detailed explanation of a Jordan chain computation for the Manhattan road network. The expensive nature of this computation is a significant drawback for many real-world applications. On the other hand, the inexact computation reveals the same eigenvector expression with computation time on the order of minutes. In addition, the AIM reduces energy dispersal among the spectral components corresponding to the zero eigenvalue.

IX. Conclusion

This paper presents an inexact transform that approximates the spectral projector-based graph Fourier transform defined in [11]. In Section II, the Agile Inexact Method (AIM) for computing the graph Fourier transform is developed for which the generalized eigenspaces are the spectral components. The generalized Parseval's identity for the inexact method is compared to that of the Jordan subspace projector method in [11]. In addition, we show that the total variation-based ordering of spectral components does not change under the inexact computation. The AIM simplifies the eigendecomposition step in real-world, sparse networks, and is associated with a fidelity-runtime trade-off. Motivated by the equivalence classes discussed in [15], we show that this transform is *agile* and describe the necessary steps for applying our method to real-world data.

We demonstrate the AIM on a graph signal based on New York City taxi trip data, which yields highly expressed locations that are compared to the ground truth. The AIM is shown to reduce the computation time for the eigendecomposition since the Manhattan road network is in the same $\mathscr{G}$-equivalence class as a graph with an eigenvector matrix that can be computed quickly. The AIM also reduces energy dispersal among the spectral components compared to oblique eigenvector projections due to the small angles between the Jordan chain vectors and their corresponding eigenvectors.

References

[1] A. Sandryhaila and J.M.F. Moura, "Discrete signal processing on graphs," *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.

[2] D. Shuman, S.K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 83–98, Apr. 2013.

[3] O. Teke and P.P. Vaidyanathan, "Extending classical multirate signal processing theory to graphs – Part I: Fundamentals," *IEEE Transactions on Signal Processing*, vol. 65, no. 2, pp. 409–422, Jan. 2017.

[4] X. Zhu and M. Rabbat, "Approximating signals supported on graphs," in *Proceedings of the 37th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Mar. 2012, pp. 3921–3924.

[5] S.K. Narang and A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *IEEE Transactions on Signal Processing*, vol. 60, no. 6, pp. 2786–2799, Jun. 2012.

[6] A.G. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Sampling of graph signals with successive local aggregations," *IEEE Transactions on Signal Processing*, vol. 64, no. 7, pp. 1832–1843, Apr. 2016.

[7] S. Segarra, A. Marques, G. Leus, and A. Ribeiro, "Reconstruction of graph signals through percolation from seeding nodes," *IEEE Transactions on Signal Processing*, vol. 64, no. 16, pp. 4363–4378, Aug. 2016.

[8] S. Chen, A. Sandryhaila, J.M.F. Moura, and J. Kovačević, "Signal recovery on graphs: Variation minimization," *IEEE Transactions on Signal Processing*, vol. 63, no. 17, pp. 4609–4624, 2015.

[9] A. Sandryhaila and J.M.F. Moura, "Big data analysis with signal processing on graphs: Representation and processing of massive data sets with irregular structure," *IEEE Signal Processing Magazine*, vol. 31, no. 5, pp. 80–90, Aug. 2014.

[10] A. Sandryhaila and J.M.F. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3042–3054, Jun. 2014.

[11] J.A. Deri and J.M.F. Moura, "Spectral projector-based graph Fourier transforms," submitted, Nov. 2016.

[12] F. Franchetti, M. Püschel, Y. Voronenko, S. Chellappa, and J. M. F. Moura, "Discrete Fourier transform on multicore," *IEEE Signal Processing Magazine*, vol. 26, no. 6, pp. 90–102, 2009.

[13] P. Lancaster and M. Tismenetsky, *The Theory of Matrices*, New York, NY, USA: Academic, 2nd edition, 1985.

[14] G.H. Golub and C.F. Van Loan, *Matrix Computations*, Baltimore, MD, USA: JHU Press, 4 edition, 2013.

[15] J.A. Deri and J.M.F. Moura, "Graph equivalence classes for spectral projector-based graph Fourier transforms," in preparation, Nov. 2016.

[16] R.A. Horn and C.R. Johnson, *Matrix Analysis*, Cambridge, U.K.: Cambridge Univ. Press, 2012.

[17] M. Püschel and J.M.F. Moura, "Algebraic signal processing theory: Foundation and 1-D time," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3572–3585, Aug. 2008.

[18] M. Püschel and J.M.F. Moura, "Algebraic signal processing theory: 1-D space," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3586–3599, Aug. 2008.

[19] M. Vetterli, J. Kovačević, and V.K. Goyal, *Foundations of Signal Processing*, Cambridge, U.K.: Cambridge Univ. Press, 2014.

[20] J.A. Deri and J.M.F. Moura, "Taxis in New York City: A network perspective," in *Proceedings of the 49th Asilomar Conference on Signals, Systems, and Computers*, Nov. 2015, pp. 1829–1833.

[21] Baruch College: Baruch Geoportal, "NYC Geodatabase," URL: https://www.baruch.cuny.edu/confluence/display/geoportal/NYC+Geodatabase, Accessed 31-Aug.-2014.

[22] L.A. Adamic and N. Glance, "The political blogosphere and the 2004 U.S. election: Divided they blog," in *Proceedings of the 3rd ACM International Workshop on Link Discovery (LinkKDD)*, 2005, pp. 36–43.

[23] W. Kahan, "Conserving confluence curbs ill-condition," Tech. Rep. AD0766916, Department of Computer Science, University of California, Berkeley, Berkeley, CA, USA, 1972.

[24] Z. Zeng, "Regularization and Matrix Computation in Numerical Polynomial Algebra," in *Approximate Commutative Algebra*, L. Robbiano and J. Abbott, Eds., Texts and Monographs in Symbolic Computation, pp. 125–162. Springer-Verlag Vienna, 2010.

[25] Z. Zeng, "Sensitivity and Computation of a Defective Eigenvalue," preprint at http://orion.neiu.edu/~zzeng/Papers/eigit.pdf, Apr. 2015.

[26] L.N. Trefethen and M. Embree, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton, NJ, USA: Princeton Univ. Press, 2005.

[27] A. Ruhe, "An algorithm for numerical determination of the structure of a general matrix," *BIT Numerical Mathematics*, vol. 10, no. 2, pp. 196–216, 1970.

[28] Axel Ruhe, "Properties of a matrix with a very ill-conditioned eigenproblem," *Numerische Mathematik*, vol. 15, no. 1, pp. 57–60, 1970.

[29] MATLAB, *version 8.5.0 (R2015a)*, The MathWorks Inc., Natick, Massachusetts, 2015.

[30] N. Guglielmi, M.L. Overton, and G.W. Stewart, "An efficient algorithm for computing the generalized null space decomposition," *SIAM Journal on Matrix Analysis & Applications*, vol. 36, no. 1, pp. 38–54, Oct. 2014.

[31] G. H. Golub and J. H. Wilkinson, "Ill-Conditioned Eigensystems and the Computation of the Jordan Canonical Form," *SIAM Review*, vol. 18, no. 4, pp. 578–619, 1976.

[32] B. Donovan and D.B. Work, "New York City Taxi Data (2010-2013)," Dataset, http://dx.doi.org/10.13012/J8PN93H8, 2014, Accessed 30-Jun.-2016.

[33] J.A. Deri, F. Franchetti, and J.M.F. Moura, "Big Data computation of taxi movement in New York City," to appear in Proceedings of the 1st IEEE Big Data Conference Workshop on Big Spatial Data, Dec. 2016.

[34] Google, "Google Map of New York, New York," https://goo.gl/maps/57U3mPQcdQ92, Accessed 31-Aug.-2014.

[35] Z. Bujavonić, *Krylov Type Methods for Large Scale Eigenvalue Computations*, Ph.D. thesis, Univ. of Zagreb, Zagreb, Croatia, 2011.

[36] J.H. Wilkinson, *The Algebraic Eigenvalue Problem*, vol. 87, Oxford, UK: Clarendon Press, 1965.

[37] J. Wilkening, "An algorithm for computing Jordan chains and inverting analytic matrix functions," *Linear Algebra and its Applications*, vol. 427, no. 1, pp. 6–25, 2007.

[38] R.T. Behrens and L.L. Scharf, "Signal processing applications of oblique projection operators," *IEEE Transactions on Signal Processing*, vol. 42, no. 6, pp. 1413–1424, 1994.